

A note about the generalisation of the C-tests

José Hernández-Orallo

DSIC, Universitat Politècnica de València, Spain

jorallo@dsic.upv.es

December 31, 2014

Abstract

In this exploratory note we ask the question of what a measure of performance for all tasks is like if we use a weighting of tasks based on a difficulty function. This difficulty function depends on the complexity of the (acceptable) policy for the task (instead of a universal distribution over tasks or an adaptive test). The resulting aggregations and decompositions are (now retrospectively) seen as the natural (and trivial) interactive generalisation of the *C*-tests.

Keywords: Intelligence evaluation, artificial intelligence, *C*-tests, algorithmic information theory, universal psychometrics, agent response curve.

1 Introduction

Since the inception of algorithmic information theory (AIT) in the 1960s, its use to construct intelligence tests was hinted by some and explicitly suggested by [7]. The first actual implementation of a test using AIT was the *C*-test [32, 16], where the goal was to find a continuation of a sequence of letters, as in some IQ tasks [55, 31], and in the spirit of Solomonoff’s inductive inference problems: “given an initial segment of a sequence, predict its continuation” (as quoted in [47, p.332]). Levin’s *Kt* complexity was used to calculate the difficulty of a sequence of letters. The performance was measured as an aggregated value over a range of difficulties:

$$I(\pi) \triangleq \sum_{h=1}^H h^e \sum_{i=1}^N \frac{1}{N} \text{Hit}(\pi, x_{i,h}) \quad (1)$$

where π is the subject, the difficulties range from $h = 1$ to H and there are N sequences $x_{i,k}$ per difficulty h . The function *hit* returns 1 if π is right with the continuation and 0 otherwise. If $e = 0$ we have that all difficulties have the same weight. The N sequences per difficulty were chosen (uniformly) randomly.

This contrasts with a more common evaluation in artificial intelligence based on average-case performance according to a probability of problems or tasks:

$$\Psi(\pi) \triangleq \sum_{\mu \in M} p(\mu) \cdot \mathbb{E}[R(\pi, \mu)] \quad (2)$$

where p is a probability distribution on the set of tasks M , and R is a result function of agent π on task μ . Actually, eq. 2 can also be combined with AIT, in a different way, by using a universal distribution [54, 47], i.e., $p(\mu) = 2^{-K(\mu)}$, where $K(\mu)$ is the Kolmogorov complexity of μ , as first chosen by [44].

The work in [44] has been considered a generalisation of [32, 16], from static sequences (predicting a continuation of a sequence correctly) to dynamic environments (maximising rewards in an MDP). In this paper we challenge this interpretation and look for a proper generalisation of [32, 16] using the notion of difficulty in the outer sum, as originally conceived and seen in eq. 1. The key idea is the realisation that for the *C*-test the task and the solution were the same thing. This meant that the difficulty was calculated as the size of the simplest programs that generates the sequence, which is both the task and the solution. Even if the complexity of the task and the solution coincide here, it is *the complexity of the solution what determines the difficulty of the problem*.

$h = 9$: a, d, g, j, ...	Answer: m
$h = 12$: a, a, z, c, y, e, x, ...	Answer: g
$h = 14$: c, a, b, d, b, c, e, c, d, ...	Answer: d

Figure 1: Several series of different difficulties 9, 12, and 14 used in the C -test [16].

However, when we move from sequences to environments or other kind of interactive tasks, the complexity of the policy that solves the task and the complexity of the environment are no longer the same. In fact, this is discussed in [24, 22]: the complexity of the environment is roughly an upper bound of the complexity of the acceptable policies (any agent that reach an acceptable performance value), but very complex environments can have very simple acceptable policies. In fact, the choice of $p(\mu) = 2^{-K(\mu)}$ has been criticised for giving too much weight to a few environments. Also, it is important to note that the invariance theorem is more meaningful for Kolmogorov Complexity than for Algorithmic Probability, as for the former it gives some stability for values of K that are not very small, but for a probability it is precisely the small cases that determine most of the distribution mass. In fact, any computable distribution can be approximated (to whatever required precision) by a universal distribution. This means that the choice of $p(\mu) = 2^{-K(\mu)}$ for Eq. 2 is actually a metadefinition, which leads to virtually any performance measure, depending on the UTM that is chosen as reference.

By decoupling the complexity of task and policy we see that we can go back to eq. 1 and work out a notion of difficulty of environments that depends on the complexity of the policy. While this may look retrospectively trivial, and the natural extension in hindsight, we need to solve and clarify some issues, and properly analyse the relation of the two different philosophies given by eq. 2 and eq. 1.

The rest of the paper is organised as follows. Section 2 discusses some previous work, introduces some notation and recovers the difficulty-based decomposition of aggregated performance. Section 3 shows what happens if difficulty functions are based on task complexity. Section 4 introduces several properties about difficulty functions and the view of difficulty as policy complexity. Section 5 discusses the choices for the difficulty-dependent probability. Section 6 briefly deals with the role of computational steps for difficulty. Section 7 closes the paper with a discussion.

2 Background

AI evaluation has been performed in many different ways [49, 14, 50, 15, 8, 9, 41, 42, 6, 53, 3, 13, 40, 43, 56, 12, 2, 48, 52] (for a recent account of AI evaluation, see [21]), but a common approach is based on averaging performance on a range of tasks, as in eq. 2.

In what follows, we will focus on the approaches that are based on AIT. As mentioned above, the first intelligence test using AIT was the so-called C -test [32, 16]. Figure 1 shows examples of sequences that appear in this test. The difficulty of each sequence was calculated as Levin’s Kt , a time-weighted version of Kolmogorov complexity K . Some preliminary experimental results showed that human performance correlated with the absolute difficulty (h) of each exercise and also with IQ test results for the same subjects. Figure 2 shows the results (taken from [32, 16]). HitRatio is defined as the inner sum of eq. 1:

$$\text{HitRatio}(\pi, h) \triangleq \sum_{i=1}^N \frac{1}{N} \text{Hit}(\pi, x_{i,h}) \quad (3)$$

An interesting observation is that by arranging problems by difficulty we see that HitRatio seems to be very small from difficulty 8 on. This makes the estimation of the measure much easier, as we only need to focus on (the area of) a small interval of difficulties. In fact, this use of difficulty is common in psychometrics.

Several generalisations of the C -test were suggested (for “cognitive agents [...] with input/output devices for a complex environment” [32] where “rewards and penalties could be used instead” [18]) or extending them for other cognitive abilities [17], but not fully developed.

AIT and reinforcement learning were finally combined in [44], where all possible environments were considered in eq. 2, instantiated with a universal distribution for p , i.e., $p(\mu) = 2^{-K(\mu)}$, with $K(\mu)$ being the Kolmogorov complexity of each environment μ . Some problems (computability, discriminating power,




Figure 2: Results obtained by humans on task of different difficulty in the C -test [16].

overweight for small environments, time, ...) were discussed with the aim of making a more applicable version of this approach by [33] and [24, secs. 3.3 and 4]. Some tests were developed [36, 35, 38, 45], using the environment class defined in [19] and some ways of aggregated rewards [20]. Despite the limited results, the experiment had quite a repercussion [39, 5, 25, 57].

While the aim of all these proposals was to measure intelligence, many interesting things can happen if AIT is applied to cognitive abilities other than intelligence, as suggested in [17] for the passive case and hinted in [24, secs. 6.5 and 7.2] for the dynamic cases, which proposes the use of different kinds of videogames as environments (two of the most recently introduced benchmarks and competitions in AI are in this direction [4, 51]). Also, some hybridisations were also proposed [27, 11, 29, 30, 34, 37], the notion of potential intelligence [26], as well as an integration with the measurement of natural systems under the umbrella of ‘universal psychometrics’ [28] and the notion of universal test [10].

2.1 Notation

We consider tests that are composed of tasks (also called environments or items) and are performed by agents (also called policies or subjects). The set of tasks is denoted by M . Its elements are usually denoted by μ . The set of agents is denoted by Π . Its elements are usually denoted by π . The length in bits of a string is denoted by $L(x)$. We can think of a proper encoding of tasks and agents as strings. Given a UTM U we define Kolmogorov complexity as $K_U(y) = \min_{x: U(x)=y} L(x)$. We will usually drop the subindex U . The expected¹ execution steps of π when performing task μ are denoted by $\mathbb{E}[S(\pi, \mu, \tau)]$ for a time limit² τ . For MDPs one possibility is to consider the (expected value of the) maximum steps taken by π for any transition [24]. If there is a transition before τ that does not halt, then $\mathbb{E}[S(\pi, \mu, \tau)]$ is infinite. The use of stochastic agents and environments makes things more complicated, as any possible transition with non-zero probability before the time limit τ that does not halt makes the expectation of S infinite. Finally, we have the expected value of the response, score or result of π in μ for a time limit τ as $\mathbb{E}[R(\pi, \mu, \tau)]$. The value of τ will be usually omitted. The R function always gives values between 0 and 1 and we assume it is always defined (a value of $R = 0$ is assumed for non-halting situations)³. We also define

¹This has to be ‘expected’ if we consider stochastic environments or agents.

²The time limit can be understood as the number of transitions, as in MDPs, or can represent a discrete or continuous time in other kind of synchronous or asynchronous modelling of tasks.

³Note that we do not go into details about how the R and S functions are calculated. One common option in reinforcement learning is $\mathbb{E}[R(\pi, \mu, \tau)] = \mathbb{E}[\sum_{i=1}^{\infty} \gamma^i \rho_i(\pi, \mu)]$, where ρ_i is the result of the reward function for iteration i . The term γ^i is a discount term with γ greater than (but usually close to) 1. The use of a discounting factor γ makes the approximation possible even in cases where the rewards are divergent. Similarly, we could also use a discount function for S as $\mathbb{E}[S(\pi, \mu, \tau)] = \mathbb{E}[\sum_{i=1}^{\infty} \gamma^i \sigma_i(\pi, \mu)]$, where σ_i gives the execution steps of transition i . Even with these choices and a discount factor, an infinite sum would make both functions very difficult to estimate, as we could have a transition with a large value of i giving a high value. For the R function this is not so critical, as it is bounded and the contribution decays with i , but for S this is problematic.




Figure 3: Two figurative agent response curves. Despite having equal areas, one is better for easier tasks while the other is better for more difficult tasks. What is more relevant (to fail in easy tasks or to excel in difficult tasks) is a matter of what we may be interested in when evaluating the systems.

$\mathbb{E}[LS(\pi, \mu, \tau)] \triangleq L(\pi) + \log \mathbb{E}[S(\pi, \mu, \tau)]$. Logarithms are always binary. Inspired by Levin’s Kt (see, e.g., [46] or [47]), we can define $Kt(\mu) = \min_{\pi} \mathbb{E}[LS(\pi, \mu)]$, which in this case depends on μ as well⁴.

2.2 Difficulty-based decomposition

It is actually in [28], where we can find a first connection between the schemas of eq. 1 and eq. 2. We adapt definition 14 in [28], which is a generalisation of eq. 2, by making the set M and the task probability p explicit as parameters.

Definition 1. *The expected average result for a task class M , a distribution p and an agent π is:*

$$\Psi(\pi, M, p) \triangleq \sum_{\mu \in M} p(\mu) \cdot \mathbb{E}[R(\pi, \mu)] \quad (4)$$

And now we use proposition 4 in [28] that decomposes it. First, we define partial results for a given difficulty h as follows:

$$\Psi_h(\pi, M, p) \triangleq \sum_{\mu \in M, h(\mu)=h} p(\mu|h) \cdot \mathbb{E}[R(\pi, \mu)] \quad (5)$$

Where h is a difficulty function $h : M \rightarrow \mathbb{R}^+ \cup 0$. Note that this parametrises the result of eq. 4 for different difficulties. For instance, for two agents π_A and π_B we might have that $\Psi_3(\pi_A) < \Psi_3(\pi_B)$ but $\Psi_7(\pi_A) > \Psi_7(\pi_B)$. If we represent $\Psi_h(\pi, M, p)$ on the y -axis versus h on the x -axis we have a so-called agent response curve, much like Fig. 2. In Fig. 3 we show an example of two agent response curves.

If we want to get a single number from an agent response curve we can aggregate performance for a range of difficulties, e.g., as follows:

Proposition 1. ([28, prop. 4]) *The expected average result $\Psi(\pi, M, p)$ can be rewritten as follows:*

$$\Psi(\pi, M, p) = \int_0^\infty p(h) \Psi_h(\pi, M, p) dh \quad (6)$$

⁴If LS is considered as a measure of *effort*, one can argue that space (i.e., memory) could also be considered, as this has been usually important for the analysis of algorithms and in cognitive science (working memory). However, execution steps are always equal or greater than working memory, so if we consider S there is no need to also include working memory, unless we are especially interested in giving relevance (and weight) to working memory.

or, in the particular case when \hbar only gives discrete values:

$$\Psi(\pi, M, p) = \sum_{h=0}^{\infty} p(h) \Psi_h(\pi, M, p) \quad (7)$$

where $p(h)$ is a probability density function for eq. 6 and a discrete probability function for eq. 7. Note that equations 4, 5 and 7 are generalisations, respectively, of equations 2, 3 and 1.

3 Difficulty as task complexity

We are now considering the decomposition of definition 1 by using two very particular choices. The first choice is that we consider $p(\mu) = 2^{-K(\mu)}$. This is exactly what has been considered in [44, 45]. The second choice is that we consider that the difficulty of an environment is $\hbar(\mu) = K(\mu)$. In this very particular case⁵, we see that \hbar is discrete (K is defined on the natural numbers).

With these two choices ($p(\mu) = 2^{-K(\mu)}$ and $\hbar = K(\mu)$), the first thing that we realise is that *for those μ such that $\hbar(\mu) = h$* we have that $p(\mu|h) = \frac{1}{N_M(h)}$, where $N_M(h) \triangleq |\{\hbar(\mu) = h\}|$. In other words, given a difficulty h , all tasks μ of that difficulty have the same K and clearly the same probability, so we have a uniform distribution. The denominator, $N_M(h)$, is (much) lower than 2^h as we are using a prefix code and some (many) sequences of length h are not self-delimited programs. The value $p(h)$ is then given by $\sum_{\mu \in M, \hbar(\mu)=h} 2^{-K(\mu)} = \sum_{\mu \in M, \hbar(\mu)=h} 2^{-h} = N_M(h)2^{-h}$. So we have:

$$\Psi(\pi, M) = \sum_{h=0}^{\infty} N_M(h)2^{-h} \sum_{\mu \in M, \hbar(\mu)=h} \frac{1}{N_M(h)} \mathbb{E}[R(\pi, \mu)]$$

Even if these are very particular (and arguable unreasonable) choices of task probability and task difficulty, this straightforward transformation is helpful to show that what we have is that all tasks of the same difficulty have the same weight (the inner sum) but the outer sum can go from a geometric distribution for difficulties (easy tasks are much more likely) if the prefix coding is a unary coding (so $N_M(h) = 1$, and $p(h) = 2^{-h}$) to much more efficient codings where $N_M(h)$ is close to 2^{h-m} where m is a constant that depends on the efficiency of the coding. For instance, for a prefix coding that uses a fixed-length word of size c (character size in bits) and a special character as end delimiter, we would have $N_M(h) = (2^c - 1)^{(\frac{h}{c}-1)}$ programs when h is divisible per c and 0 otherwise. In this case $p(h) = N_M(h)2^{-h}$ can be approximated by $\frac{1}{c \cdot \nu} b^{-h}$, where $b = 1 + 2^{-c}$ and ν is a constant normalisation term such that $\sum w(h) \leq 1$. This is what we see on Figure 4 (left) with $c = 4$ (so $b = 1.0625$). However, we see that $w(h)$ can have a base as close to 1 as we like, so the probability would look uniform for small values of h , as Figure 4 (right). Nonetheless, it still is a geometrical distribution (otherwise the sum would diverge, which is not possible as we assume that R is bounded as for equation 4).

All this highlights how important the prefix coding is for the understanding of these distributions. In fact, it is not only that the definition depends strongly on the choice of the reference UTM used for K , which determines the probability of each task $2^{-K(\mu)}$, as argued elsewhere ([33] and [24, secs. 3.3 and 4]), but also that the prefix coding is highly relevant as well if we use $K(\mu)$ as difficulty.

Apart from the differences that appear because of the coding, there is a major concern about this: it is now crystal clear that easier problems would have more weight than difficult ones. However, our intuition when evaluating a subject on a range of difficulties would be to give more relevance to more difficult problems or, at most, to give the same relevance to all difficulties. In the latter case, a subject being able to score well in all tasks of difficulty $2h$ would double the result than another subject being to score well in all tasks of difficulty h .

4 Difficulty functions

Before setting an appropriate measure of difficulty based on the policy, in this section we will analyse which properties a difficulty function may have.

⁵In [28, prop. 3], it is shown that with these two choices the average difficulty diverges. This does not mean that eq. 7 diverges, which is what we work out here.




Figure 4: Left: $p(h) = N_M(h)2^{-h}$ with a character coding with one end delimiter with word length $c = 4$. Right: $p(h)$ with word length $c = 12$. In blue the true values and in red the approximation $\frac{1}{c \cdot \nu} b^{-h}$, where $b = 1 + 2^{-c}$ ($b = 1.0625$ on the left and $b = 1.000244$ on the right). The distribution on the right looks almost uniform but it still geometric.

4.1 Detaching $p(h)$ from $p(\mu|h)$

The decomposition in previous section suggests that we could try to fix a proper measure of difficulty first and then think about a meaningful distribution $p(h)$. Once this is settled, we could try to find a distribution for all environments of that difficulty $p(\mu|h)$. In other words, once we determine how relevant a difficulty is we ask which tasks to take for that difficulty. This is the spirit of the C -test [32, 16] as seen in eq. 1. In fact, we perhaps we do not need a $p(h)$ that decays dramatically, as it is expectable to see performance to decrease for increasing difficulty, as in Figure 2.

To distinguish $p(h)$ and $p(\mu|h)$ we will denote the former with w and the latter with p_M . We will use any distribution or even a measure (not summing up to one, for reasons that we will see later on) as a subscript for Ψ . For instance, we will use the following notation $\Psi_{\mathcal{U}(h_{min}, h_{max})}(\pi, M, p_M)$, where $\mathcal{U}(a, b)$ represents a uniform distribution between a and b . For instance, we can have two agents π_A and π_B such that $\Psi_{\mathcal{U}(1,10)}(\pi_A) > \Psi_{\mathcal{U}(1,10)}(\pi_B)$ but $\Psi_{\mathcal{U}(11,20)}(\pi_A) < \Psi_{\mathcal{U}(11,20)}(\pi_B)$. We will use the notation $\Psi_{\oplus}(\pi, M, p_M)$ when $w(h) = 1$ (note that this is not the uniform distribution for discrete h), which means that the partial aggregations for each difficulty are just added. In other words, $\Psi_{\oplus}(\pi, M, p_M) \triangleq \sum_{h=0}^{\infty} \Psi_h(\pi, M, p_M)$ for discrete difficulty functions and $\Psi_{\oplus}(\pi, M, p_M) \triangleq \int_0^{\infty} \Psi_h(\pi, M, p_M) dh$ for continuous difficulty functions. We will explore whether this (area under the agent response curve) is bounded.

Figure 5 shows approach A, which has already been mentioned, while approaches B and C will be seen in sections 5.1 and 5.2 respectively.

4.2 Acceptable policies

When we aggregate environments with different scales on R and different difficulties, we may have that an agent focusses on a few environments with high difficulty while another focusses on many more environments with small responses. Agent response curves in [28], which are inspired by item response curves in psychometrics (but inverting the view between agents and items), allow us to see how each agent performs for different degrees of difficulty. Looking at Figure 2 and similar agent response curves in psychometrics, we see that the notion of difficulty must be linked to R , i.e., how well the agents perform, and not about the complexity of the task, as in the previous section.

One first idea of a difficulty function of an environment is the expected response for a random agent




Figure 5: Three approaches to aggregate the results for a set of tasks. Top (A) shows the classical approach of choosing a probability for the task, according to the properties of the task. Middle (B) shows the approach where we arrange tasks by difficulty, and the notion of difficulty is derived from the properties of the policy. Bottom (C) shows a variation of B where we derive acceptable policies for a given difficulty and then generate tasks for each policy. Between square brackets some choices we examine in this paper.




Figure 6: Left: an illustrative distribution of responses of a population of agents for a single environment, assuming a finite number of agents. A random agent π_{rand} has expected response of 0.46. Right: the distribution of responses ordered by $K(\pi)$ on the x -axis.

π_{rand} , i.e., $\bar{h}(\mu) \triangleq \mathbb{E}[R(\pi_{rand}, \mu)]$. However, this result is not very meaningful, but just sets some kind of baseline, which will be between 0 and 1. This is shown in Figure 6 (left) and has the advantage of being applicable to infinite populations of agents Π without assuming any distribution of agents. A related idea would be to calculate the expected response for a population of agents, using some distribution, e.g., $\bar{h}(\mu) \triangleq \sum_{\pi} p_{\Pi}(\pi) \mathbb{E}[R(\pi, \mu)]$. As agents are programs, we may have the temptation of using $p_{\Pi}(\pi) = 2^{-K(\pi)}$. If we look at this for an illustrative example in Figure 6 (right), we see that this would mean that basically the few results on the leftmost part of the plot will dominate, i.e., the result of the shortest program.

As we want a short policy with a good result, we can think about a ratio between complexity and response of the policy, something like $\bar{h}(\mu) \triangleq \min_{\pi} \frac{L(\pi)}{\mathbb{E}[R(\pi, \mu)]}$. The unit for difficulty would be bits per reward unit. It is not clear why doubling responses should take double bits (perhaps logarithms could be used instead). In the example in Figure 6 (right), we see that the leftmost policy ($L(\pi) = 1$) would have $h = \frac{1}{0.42} = 2.38$, and the best with $L(\pi) = 2$ would have $h = \frac{2}{0.78} = 2.56$. In fact, no other policy could get lower than in 2.38. Choosing any other reasonable function instead of a ratio such that we can go further right on the plot would be rather arbitrary. Another option is what is done in [22], as $\bar{h}(\mu) \triangleq \min_{\pi: \mathbb{E}[R(\pi, \mu)] = R_{max}(\mu)} L(\pi)$ where $R_{max}(\mu) = \max_{\pi} \mathbb{E}[R(\pi, \mu)]$. However, R_{max} may be hard to calculate and even if it can be effectively calculated, any minor mistake or inefficiency in a very good agent will prevent the agent from reaching the

optimal result, leading to a notion of difficulty linked to the complexity of the ‘perfect’ policy. In [22], a ‘tolerance value’ is considered and, instead of one policy, difficulty is linked to the probability of finding a policy under this tolerance by using different search approaches.

Here we are going to consider a similar approach. We are going to consider this tolerance ϵ of acceptability.

$$\mathbb{A}^{[\epsilon]}(\pi, \mu) \triangleq \mathbf{1}(\mathbb{E}[R(\pi, \mu)] \geq 1 - \epsilon) \quad (8)$$

This returns 1 if the expected response is above $1 - \epsilon$ and 0 otherwise. If $\mathbb{A}^{[\epsilon]}(\pi, \mu) = 1$ we say that π is ϵ -acceptable. With this, we binarise responses. One can argue that we could have just defined a binary R , but it is important to clarify that it is not the same to have tolerance for each single R (or a binarised R) than to have a tolerance for the expected value $\mathbb{E}[R]$. The tolerance on the expected value allows the agent to have variability in their results (e.g., stochastic agents) provided the expected value is higher than the tolerance. Finally, even if we will be very explicit about the value of ϵ , and changing it will change the difficulty value of any environment, it is important to say that this value is not so relevant. The reason is that for any environment we can build any other environment where the responses are transformed by any function. In fact, we could actually consider one fixed threshold, such as 0.5, always.

And now we can just define a new version of eq. 5 using this new function:

$$\Psi_h^{[\epsilon]}(\pi, M, p_M) \triangleq \sum_{\mu \in M, \hbar(\mu)=h} p_M(\mu|h) \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \quad (9)$$

We can just rewrite equations 6 and 7 accordingly:

$$\Psi_w^{[\epsilon]}(\pi, M, p_M) = \int_0^\infty w(h) \Psi_h^{[\epsilon]}(\pi, M, p_M) dh \quad (10)$$

$$\Psi_w^{[\epsilon]}(\pi, M, p_M) = \sum_{h=0}^\infty w(h) \Psi_h^{[\epsilon]}(\pi, M, p_M) \quad (11)$$

4.3 Properties

Given the above, we are now ready for a few properties about difficulty functions.

Definition 2. A difficulty function \hbar is strongly bounded in M if for every π there is a difficulty h such that for every $\mu \in M : \hbar(\mu) \geq h$ we have $\mathbb{A}^{[\epsilon]}(\pi, \mu) = 0$.

A weaker version of the above property would be as follows:

Definition 3. A difficulty function \hbar is weakly bounded in M if for every π there is a difficulty h such that for every $h' \geq h$ we have $\Psi_{h'}^{[\epsilon]}(\pi, M, p_M) = 0$.

Clearly, strongly boundedness implies weakly boundeness, but not vice versa. There can be cases that are not strongly bounded but are weakly bounded. For instance, if we just choose $p_M(\mu|h) = 0$ if $h > h'$ for a given h' , which is not actually because of a difficulty function, so this is show that weakly boundness is rather a property of the evaluation setting and not only about the difficulty function.

Proposition 2. If a countable difficulty function \hbar is weakly bounded by R then for every possible agent π we have that:

$$\Psi_{\oplus}^{[\epsilon]}(\pi, M, p_M) = \sum_{h=0}^\infty \Psi_h^{[\epsilon]}(\pi, M, p_M) \quad (12)$$

is finite.

Proof. As it is weakly bounded then there is an h such that for every $h' \geq h$ we have $\Psi_{h'}^{[\epsilon]}(\pi, M, p_M) = 0$. Consequently, $\sum_{h=0}^\infty \Psi_h^{[\epsilon]}(\pi, M, p_M) = \sum_{h=0}^{h'} \Psi_h^{[\epsilon]}(\pi, M, p_M)$. As \hbar is countable, and Ψ_h is bounded, then eq. 12 is finite. \square

The above is a modification of eq. 11 where the probability $p(h)$ has been removed. When the above proposition holds we have that the area under the agent response curve is finite. We can obtain a similar result for the continuous case (eq. 10), using an integral instead.




Figure 7: Left: an illustrative distribution of responses of a population of agents for a single environment. If we set the threshold at $0.9 = 1 - \epsilon$, the simplest policy above this threshold is of ‘complexity’ $h = 12$. Right: an illustrative distribution of the result of Ψ_h considering a population of environments for a single agent, an *agent response curve*. There is no task π of difficulty above 80 for which $\mathbb{E}[R(\pi, \mu)] \geq 1 - \epsilon$, i.e., there is no task for which π is ϵ -acceptable, so Ψ_h is 0 from 80 on. If we were using the definition of \bar{h} as for Eq. 13, this 80 would be $L(\pi)$. Also note on the right plot that all ‘heaven’ tasks (good results independently of what the agent does) are at $h = 1$, while all ‘hell’ tasks (bad response independently of what the agent does) are at $h = \infty$.

4.4 Difficulty as policy complexity

Now we are ready to ask what happens if we choose the difficulty function in terms of ϵ -acceptability, i.e.:

$$\bar{h}^{[\epsilon]}(\mu) \triangleq \min\{L(\pi) : \mathbb{E}[R(\pi, \mu)] \geq 1 - \epsilon\} = \min\{L(\pi) : \mathbb{A}^{[\epsilon]}(\pi, \mu) = 1\} \quad (13)$$

We can say a few words about the cases where a truly random agent gives an acceptable policy for an environment. If this is the case, we intuitively consider the environment easy. So, in terms, of L , we consider truly random agents to be simple, which is more reasonable than considering them of infinite difficulty, and goes well with our consideration of stochastic agents and environments.

Figure 7 (left) shows the distribution of response according to $L(\pi)$, but setting $\epsilon = 0.9$. We see that the simplest ϵ -acceptable policy has $L = 12$.

With the difficulty function in eq. 13 we have:

Proposition 3. *The difficulty function $\bar{h}^{[\epsilon]}$ in eq. 13 is strongly bounded.*

Proof. For every policy π , if a task μ has a difficulty $\bar{h}^{[\epsilon]}(\mu) > L(\pi)$, it means that π is not ϵ -acceptable, because otherwise the difficulty would be $L(\pi)$ and not h . Consequently, $\mathbb{A}^{[\epsilon]}(\pi, \mu) = 0$ for all μ of difficulty $\bar{h}^{[\epsilon]}(\mu) > L(\pi)$. It is sufficient to take $h > L(\pi)$ for every π to see that \bar{h} is strongly bounded. \square

This is what we see in Fig. 7 (right), where $L(\pi) = 80$. With $\bar{h}^{[\epsilon]}$ in eq. 13, we can ensure that the values are going to be 0 from $h = 80$ on.

This may not be the case for other difficulty functions. We can imagine a situation where the curve never converges to zero. For instance, if the difficulty function is decoupled from resources (length and/or steps) of the acceptable policies or we do not use the notion of ϵ -acceptability then we cannot avoid that a very simple policy could eventually score well in a problem with very high difficulty. This would be counter-intuitive, as if there is a simple policy for a difficult problem, the latter should not be considered difficult any more.

Finally, it is an interesting question to determine what happens with this definition of difficulty when we consider not only the shortest policy but all the policies that are ϵ -acceptable. Below, we just include one

possibility, where we weight all the acceptable policies by some extreme distribution so that we can keep the original properties.

$$h^{[\epsilon]}(\mu) \triangleq -0.5 \log \sum_{\pi} 2^{-2L(\pi)} \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \quad (14)$$

The above expression is tweaked so that in the case there is only one acceptable policy, the function gives the same result as eq. 13 (if the sum only has one the case it reduces to $L(\pi)$). In a way, this is a more robust version of difficulty than eq. 13. Still, with the strong weighting ($2^{-2L(\pi)}$) we have used, we can see that strongly boundedness also holds for this version:

Proposition 4. *The difficulty function $h^{[\epsilon]}$ in eq. 14 is strongly bounded.*

Proof. Assume $k = L(\pi)$, the shortest acceptable policy for μ . Let us consider one extreme (worst) case such that all agents of the same or more length are also acceptable. Since there are not more than 2^n programs of length n (actually fewer if it is a prefix coding), we have $h^{[\epsilon]} = -0.5 \log \sum_{n \geq k} 2^n 2^{-2n} = -0.5 \log \sum_{n \geq k} 2^{-n} = -0.5 \log 2^{-k+1} = 0.5(k-1)$. For the other extreme (best) case, which is when we only have one acceptable policy for μ , we have $h^{[\epsilon]} = -0.5 \log 2^{-2k} = k$. As this latter value is larger, we apply the same argument than for proposition 3 with $h > L(\pi)$. \square

As the above difficulty function is more complex, in what follows, we will work with the difficulty function $h^{[\epsilon]}$ in eq. 13.

5 Difficulty-conditional task probabilities

In the previous sections we have focussed on $w(h)$ and whether it is necessary or not. We have seen difficulty functions where just aggregating Ψ_h without $w(h)$ (or $w(h) = 1$) leads to a $\Psi_{\oplus}(\pi, M, p_M)$ that is bounded (proposition 2). The question now is how to choose the conditional probability $p_M(\mu|h)$. In the C -test, eq. 1, this was chosen as a uniform distribution. However, this is not possible in an interactive scenario if we consider all possible tasks, as the number of tasks for which there is an acceptable policy π of $L(\pi) = n$ can be infinite. Even if we cannot set a uniform distribution, we want a choice of $p_M(\mu|h)$ that keeps the task diversity (unless there is any special bias to choose the tasks).

5.1 Task probability depends on difficulty

The first thing we can do is to assume $p(\mu|h)$ in eq. 9 as $p(\mu|h) = \frac{2^{-K(\mu)}}{\nu(h)}$ if $h^{[\epsilon]}(\mu) = h$ and 0 otherwise, where $\nu(h)$ is a normalisation term to make the mass of the distribution equal to 1, which can be formulated as $\nu(h) = \sum_{\mu: h^{[\epsilon]}(\mu)=h} 2^{-K(\mu)}$.

And now we have:

$$\begin{aligned} \Psi_h^{[\epsilon]}(\pi, M, p_M) &= \sum_{\mu \in M, h^{[\epsilon]}(\mu)=h} p_M(\mu|h) \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \\ &= \frac{1}{\nu(h)} \sum_{\mu \in M, h^{[\epsilon]}(\mu)=h} 2^{-K(\mu)} \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \end{aligned}$$

From here, we can plug it into eq. 11 for the discrete case:

$$\Psi_w^{[\epsilon]}(\pi, M, p_M) = \sum_{h=0}^{\infty} w(h) \frac{1}{\nu(h)} \sum_{\mu \in M, h^{[\epsilon]}(\mu)=h} 2^{-K(\mu)} \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \quad (15)$$

Note that the above is going to be bounded independently of the difficulty function if w is a probability distribution. Also notice that $\frac{1}{\nu(h)}$ is on the outer sum, and that $\nu(h)$ is lower than 1, so the normalisation term is actually greater than 1.

And if we use any of the difficulty functions in equations 13 or 14 we can choose $w(h) = 1$ and $\Psi_{\oplus}^{[\epsilon]}(\pi, M, p_M)$ is bounded.

5.2 Task probability depends on the policy probability

One of things of the use of equation 13 is that the number of acceptable policies per difficulty is finite. This is what happened in the C -test and that is the reason why a uniform distribution could be used for the inner sum. We could try to decompose the inner sum by using the policy and get the probability of the task given the policy.

We first need to define this set:

$$\begin{aligned} Pairs(M, \Pi, \epsilon) &\triangleq \{ \langle \mu, \pi \rangle : \mu \in M, \pi \in \Pi, \text{ and } \pi \text{ is an } \epsilon\text{-acceptable policy for } \mu \} \\ &= \{ \langle \mu, \pi \rangle : \mu \in M, \pi \in \Pi, \text{ and } \mathbb{A}^{[\epsilon]}(\pi, \mu) = 1 \} \end{aligned}$$

Note that one task can have many acceptable policies (and perhaps none) and one agent can be an acceptable policy of many tasks (and perhaps none). Now we can have an alternative expression of difficulty equivalent to eq. 13 as follows:

$$h^{[\epsilon]}(\mu) = \min_{\langle \mu, \pi \rangle \in Pairs(M, \Pi, \epsilon)} L(\pi)$$

We want to set $p_M(\mu|h)$ by decomposing it into $p_M(\mu|\pi)$ and $p_M(\pi|h)$.

$$\begin{aligned} \Psi_h^{[\epsilon]}(\pi, M, p_M) &= \sum_{\mu \in M, h^{[\epsilon]}(\mu)=h} p_M(\mu|h) \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \\ &= \sum_{\pi': K(\pi')=h} p_M(\pi'|h) \sum_{\mu \in M: \langle \mu, \pi' \rangle \in Pairs(M, \Pi, \epsilon)} p_M(\mu|\pi') \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \end{aligned}$$

If we use the difficulty function in equation 13 we can assume that $p_M(\pi'|h)$ is uniform. If $N_\Pi(h)$ denotes the number of programs π' for which $L(\pi') = h$, we can assume $p_M(\pi'|h) = \frac{1}{N(h)}$. So we can rewrite

$$\Psi_h^{[\epsilon]}(\pi, M, p_M) = \sum_{\pi': K(\pi')=h} \frac{1}{N_\Pi(h)} \sum_{\mu \in M: \langle \mu, \pi' \rangle \in Pairs(M, \Pi, \epsilon)} p_M(\mu|\pi') \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu)$$

Now we may have the temptation to define $p_M(\mu|\pi')$ with a conditional universal distribution, i.e., $2^{-K(\mu|\pi')}$. There is a conundrum here, as we would have that those environments that can be easily described from the policy would be more likely. Even if K is not symmetric there are strong relations in both directions such that the notion of difficulty would be completely mangled. Instead, it seems more reasonable to choose the distribution as $p_M(\mu|\pi') = \frac{2^{-K(\mu)}}{\nu(\pi')}$ if $\langle \mu, \pi' \rangle \in Pairs(M, \Pi)$ and 0 otherwise, where $\nu(\pi')$ is a normalisation term, which can be calculated as $\nu(\pi') = \sum_{\langle \mu, \pi' \rangle \in Pairs(M, \Pi)} 2^{-K(\mu)}$.

We can integrate the above into eq. 11 again.

$$\begin{aligned} \Psi_w^{[\epsilon]}(\pi, M, p_M) &= \sum_{h=0}^{\infty} w(h) \sum_{\pi: K(\pi)=h} \frac{1}{N_\Pi(h)} \sum_{\mu \in M: \langle \mu, \pi \rangle \in Pairs(M, \Pi, \epsilon)} \frac{2^{-K(\mu)}}{\nu(\pi)} \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \\ &= \sum_{h=0}^{\infty} \frac{w(h)}{N_\Pi(h)} \sum_{\pi: K(\pi)=h} \frac{1}{\nu(\pi)} \sum_{\mu \in M: \langle \mu, \pi \rangle \in Pairs(M, \Pi, \epsilon)} 2^{-K(\mu)} \cdot \mathbb{A}^{[\epsilon]}(\pi, \mu) \end{aligned} \quad (16)$$

The interpretation would be as follows: for each difficulty value we aggregate all the acceptable policies with size equal to that difficulty uniformly and for each of these policies all the environments where each policy is acceptable with a universal distribution. This extra complication with respect to eq. 15 can only be justified if we generate environments and agents and we check them as we populate $Pairs$, as a way of constructing a test more easily. Once a sample of $Pairs$ is generated we may want to see how to arrange everything and give appropriately weights to each case.

Apart from this, eq. 15 seems more natural and simpler than eq. 16 as an extension of eq. 1.

6 Using computational steps

As we mentioned in the introduction, the C -test [32, 16] used Levin’s Kt instead of K . Apart from being computable, Kt is related to Levin’s search. This makes its choice much more appropriate for a measure of difficulty. However, when working with interactive tasks and with stochastic tasks and agents, we have that the number of computational steps must be calculated as $\mathbb{E}[S(\pi, \mu)]$. For simplicity, up to this point, we have considered just K . Now we are briefly exploring the inclusion of the computational steps. In section 2.1 we defined LS . Now we define a version that accounts for the tolerance ϵ as follows:

$$LS^{[\epsilon]}(\pi, \mu) \triangleq \mathbb{E}[LS(\pi, \mu)] \text{ if } \mathbb{A}^{[\epsilon]}(\pi, \mu) = 1 \text{ and } \infty \text{ otherwise}$$

and we define a new difficulty function that considers computational steps:

$$h^{[\epsilon]}(\mu) \triangleq \min_{\pi} LS^{[\epsilon]}(\pi, \mu)$$

This difficulty function is not bounded, as LS depends on μ , and we can always find a very short policy that takes an enormous amount of steps for a task with very high difficulty. This is an acceptable policy, but does not reduce the difficulty of the task, so it can always score non-zero beyond any limit.

This means that for this difficulty function we would need to use equation eq. 11 or eq. 10 with an appropriate $w(h)$ (e.g., a small decay or a uniform interval of difficulties we are interested in).

If the testing procedure established a limit on the number of steps (total or per transition) we would have this new difficulty function would be strongly bounded. Alternatively, we could reconsider the inclusion the computational steps in the notion of acceptability. In this case, the approach in section 5.2 could not be used, as the probability of π given h would also depend on μ .

7 Discussion

We have gone from eq. 1 taken from C -test to eq. 11 when using discrete difficulty functions or to eq. 10 when using continuous difficulty functions. We have seen that difficulties allow for a more detailed analysis of what happens for a given agent, depending on whether it succeeds at easy or difficult tasks. For some difficulty functions (e.g., those for which proposition 2 is true), we do not even need to determine the weight for each difficulty and just calculate the area, as an aggregated performance for all difficulties. Actually, given the characteristics of most of the difficulty functions we have seen, setting a maximum difficulty could be enough, for instance, $h = 200$, more than convenient, as it is really challenging to look into a space of policies of 200 bits (without prior knowledge). So, a finite range of difficulties would be enough for the evaluation of feasible agents.

The important thing is that now we do not need to set an a priori distribution for all tasks $p(\mu)$, but just a conditional distribution $p(\mu|h)$. Note that if we set a high h we have the freedom to find the simplest task that creates that difficulty. In fact, this is desirable, as when we try to create a difficult game with a small number of rules. Actually, the choice of $p(\mu|h)$ as a universal distribution still depends on the reference machine and can set most of the probability mass on smaller tasks, but as it is conditional on h , all trivial, dead or simply meaningless tasks have usually very extreme values of h (very low or infinite). That means that there is a range of *interesting* difficulties, discarding very small values of h and very large values of h . Figure 2 is a nice example of this, where only difficulties between 1 and 8 were used, and we see also that $h = 1$ and $h > 7$ are not really very discriminating. The bulk of the testing effort must be performed in this range.

Now, let us spare a comment about the practical feasibility of generating a test that is conceived in terms of difficulty, which is determined from the length of the policy. It is pertinent to quote a piece from [45]: “Another important difference in our work is that we have directly sampled from program space. This is analogous to the conventional construction of the Solomonoff prior, which samples random bit sequences and treats them as programs. With this approach all programs that compute some environment count towards the environment’s effective complexity, not just the shortest, though the shortest clearly has the largest impact. This makes AIQ very efficient in practice since we can just run sampled programs directly, avoiding the need to have to compute complexity values through techniques such as brute force program search. For

example, to compute the complexity of a 15 symbol program, the C-test required the execution of over 2 trillion programs. For longer programs, such as many that we have used in our experiments, this would be completely intractable. One disadvantage of our approach, however, is that we never know the complexity of any given environment; instead we know just the length of one particular program that computes it.”. While I agree that the calculation of difficulties may be hard, a couple of things must be said. First, yet again the previous quote talks about the *complexity of an environment*, which is based on its description (or on its Solomonoff prior), but not on its policies. Second, it is always preferable to devote time to design good tests that lead to a few tasks that are discriminating and whose difficulty we are sure of than to have a test that can be obtained more efficiently but the agent will take thousands of useless tasks. In other words, we want tests that are practical when they are applied. We can devote much more resources (and time) for the offline generation of appropriate tasks and the calculation of their difficulty. This is actually what real measurement disciplines do, such as psychometrics, which devote most of the effort to get good tasks and do not hesitate to discard those tasks whose difficulty is dubious or are not discriminating. Effort in the assessment and classification of exercises pays off in more efficient evaluations.

Finally, there is an important question about the choice of a meaningful difficulty function linked to the effort required to find an acceptable policy: what if difficulty depends on the verification of the policy? This has got attention for problems where the policy must be accompanied by a verification, proof or explanation [18, 1]. In other words, the question is whether we can extend Levin’s search by considering the verification of stochastic problems and derive a difficulty function from it. This is being addressed in a separate note [23].

References

- [1] T. Alpcan, T. Everitt, and M. Hutter. Can we measure the difficulty of an optimization problem? *IEEE Information Theory Workshop (ITW)*, 2014. 13
- [2] J. Anderson, J. Baltes, and C. T. Cheng. Robotics competitions as benchmarks for AI research. *The Knowledge Engineering Review*, 26(01):11–17, 2011. 2
- [3] D. Baldwin and S. B. Yadav. The process of research investigations in artificial intelligence - a unified view. *Systems, Man and Cybernetics, IEEE Transactions on*, 25(5):852–861, 1995. 2
- [4] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 06 2013. 3
- [5] C. Biever. Ultimate IQ: one test to rule them all. *New Scientist*, 211(2829,10 September 2011):42–45, 2011. 3
- [6] B. G. Buchanan. *Artificial intelligence as an experimental science*. Springer, 1988. 2
- [7] G. J. Chaitin. Gödel’s theorem and information. *International Journal of Theoretical Physics*, 21(12):941–954, 1982. 1
- [8] P. R. Cohen and A. E. Howe. How evaluation guides AI research: The message still counts more than the medium. *AI Magazine*, 9(4):35, 1988. 2
- [9] K. S. Decker, E. H. Durfee, and V. R. Lesser. Evaluating research in cooperative distributed problem solving. *Distributed Artificial Intelligence*, 2:487–519, 1989. 2
- [10] D. L. Dowe and J. Hernández-Orallo. How universal can an intelligence test be? *Adaptive Behavior*, 22(1):51–69, 2014. 3
- [11] D. L. Dowe, J. Hernández-Orallo, and P. K. Das. Compression and intelligence: social environments and communication. In J. Schmidhuber, K.R. Thórisson, and M. Looks, editors, *Artificial General Intelligence*, volume 6830, pages 204–211. LNAI series, Springer, 2011. 3
- [12] C. Drummond and N. Japkowicz. Warning: statistical benchmarking is addictive. Kicking the habit in machine learning. *Journal of Experimental & Theoretical Artificial Intelligence*, 22(1):67–80, 2010. 2
- [13] E. Falkenauer. On method overfitting. *Journal of Heuristics*, 4(3):281–287, 1998. 2

- [14] J. Gaschnig, P. Klahr, H. Pople, E. Shortliffe, and A. Terry. Evaluation of expert systems: Issues and case studies. *Building expert systems*, 1:241–278, 1983. 2
- [15] J. R. Geissman and R. D. Schultz. Verification & validation. *AI Expert*, 3(2):26–33, 1988. 2
- [16] J. Hernández-Orallo. Beyond the Turing Test. *J. Logic, Language & Information*, 9(4):447–466, 2000. 1, 2, 3, 6, 12
- [17] J. Hernández-Orallo. Computational measures of information gain and reinforcement in inference processes. *AI Communications*, 13(1):49–50, 2000. 2, 3
- [18] J. Hernández-Orallo. On the computational measurement of intelligence factors. In A. Meystel, editor, *Performance metrics for intelligent systems workshop*, pages 1–8. National Institute of Standards and Technology, Gaithersburg, MD, 2000. 2, 13
- [19] J. Hernández-Orallo. A (hopefully) non-biased universal environment class for measuring intelligence of biological and artificial systems. In M. Hutter et al., editor, *Artificial General Intelligence, 3rd Intl Conf*, pages 182–183. Atlantis Press, Extended report at <http://users.dsic.upv.es/proy/anynt/unbiased.pdf>, 2010. 3
- [20] J. Hernández-Orallo. On evaluating agent performance in a fixed period of time. In M. Hutter et al., editor, *Artificial General Intelligence, 3rd Intl Conf*, pages 25–30. Atlantis Press, 2010. 3
- [21] J. Hernández-Orallo. AI evaluation: past, present and future. *arXiv preprint arXiv:1408.6908*, 2014. 2
- [22] J. Hernández-Orallo. On environment difficulty and discriminating power. *Autonomous Agents and Multi-Agent Systems*, pages 1–53, 2014. 2, 7, 8
- [23] J. Hernández-Orallo. Universal psychometrics tasks: difficulty, composition and decomposition. *arXiv preprint*, 2015. 13
- [24] J. Hernández-Orallo and D. L. Dowe. Measuring universal intelligence: Towards an anytime intelligence test. *Artificial Intelligence*, 174(18):1508–1539, 2010. 2, 3, 5
- [25] J. Hernández-Orallo and D. L. Dowe. Mammals, machines and mind games. Who’s the smartest? *The Conversation*, <http://theconversation.edu.au/articles/mammals-machines-and-mind-games-whos-the-smartest-1125>, 2011. 3
- [26] J. Hernández-Orallo and D. L. Dowe. On potential cognitive abilities in the machine kingdom. *Minds and Machines*, 23:179–210, 2013. 3
- [27] J. Hernández-Orallo, D. L. Dowe, S. España-Cubillo, M. V. Hernández-Lloreda, and J. Insa-Cabrera. On more realistic environment distributions for defining, evaluating and developing intelligence. In J. Schmidhuber, K.R. Thórisson, and M. Looks, editors, *Artificial General Intelligence*, volume 6830, pages 82–91. LNAI, Springer, 2011. 3
- [28] J. Hernández-Orallo, D. L. Dowe, and M. V. Hernández-Lloreda. Universal psychometrics: Measuring cognitive abilities in the machine kingdom. *Cognitive Systems Research*, 27:5074, 2014. 3, 4, 5, 6
- [29] J. Hernández-Orallo, J. Insa-Cabrera, D. L. Dowe, and B. Hibbard. Turing machines and recursive Turing Tests. In V. Muller and A. Ayesh, editors, *AISB/IACAP 2012 Symposium “Revisiting Turing and his Test”*, pages 28–33. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour, 2012. 3
- [30] J. Hernández-Orallo, J. Insa-Cabrera, D. L. Dowe, and B. Hibbard. Turing Tests with Turing machines. In Andrei Voronkov, editor, *Turing-100*, volume 10, pages 140–156. EPiC Series, 2012. 3
- [31] J. Hernández-Orallo, F. Martínez-Plumed, U. Schmid, M. Siebers, and D. L. Dowe. Computer models solving human intelligence test problems: progress and implications. *submitted*, 2014. 1
- [32] J. Hernández-Orallo and N. Minaya-Collado. A formal definition of intelligence based on an intensional variant of Kolmogorov complexity. In *Proc. Intl Symposium of Engineering of Intelligent Systems (EIS’98)*, pages 146–163. ICSC Press, 1998. 1, 2, 6, 12
- [33] B. Hibbard. Bias and no free lunch in formal measures of intelligence. *Journal of Artificial General Intelligence*, 1(1):54–61, 2009. 3, 5

- [34] J. Insa-Cabrera, J. L. Benacloch-Ayuso, and J. Hernández-Orallo. On measuring social intelligence: Experiments on competition and cooperation. In J. Bach, B. Goertzel, and M. Iklé, editors, *AGI*, volume 7716 of *Lecture Notes in Computer Science*, pages 126–135. Springer, 2012. 3
- [35] J. Insa-Cabrera, D. L. Dowe, S. España-Cubillo, M. V. Hernández-Lloreda, and J. Hernández-Orallo. Comparing humans and AI agents. In J. Schmidhuber, K.R. Thórisson, and M. Looks, editors, *Artificial General Intelligence*, volume 6830, pages 122–132. LNAI, Springer, 2011. 3
- [36] J. Insa-Cabrera, D. L. Dowe, and J. Hernández-Orallo. Evaluating a reinforcement learning algorithm with a general intelligence test. In J.A. Moreno J.A. Lozano, J.A. Gamez, editor, *Current Topics in Artificial Intelligence. CAEPIA 2011*. LNAI Series 7023, Springer, 2011. 3
- [37] J. Insa-Cabrera and J. Hernández-Orallo. Definition and properties to assess multi-agent environments as social intelligence tests. *arXiv preprint*, <http://arxiv.org/abs/1408.6350>, 2014. 3
- [38] J. Insa-Cabrera, J. Hernández-Orallo, D.L. Dowe, S. España, and M.V. Hernández-Lloreda. The ANYNT project intelligence test : Lambda - one. In V. Muller and A. Ayesh, editors, *AISB/IACAP 2012 Symposium “Revisiting Turing and his Test”*, pages 20–27. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour, 2012. 3
- [39] K. Kleiner. Who are you calling bird-brained? An attempt is being made to devise a universal intelligence test. *The Economist*, 398(8723, 5 March 2011):82, 2011. 3
- [40] J. Langford. Clever methods of overfitting. *Machine Learning (Theory)*, <http://hunch.net/?p=22>, 2005. 2
- [41] P. Langley. Research papers in machine learning. *Machine Learning*, 2(3):195–198, 1987. 2
- [42] P. Langley. The changing science of machine learning. *Machine Learning*, 82(3):275–279, 2011. 2
- [43] S. Legg and M. Hutter. Tests of machine intelligence. In Max Lungarella, Fumiya Iida, Josh Bongard, and Rolf Pfeifer, editors, *50 Years of Artificial Intelligence*, volume 4850 of *Lecture Notes in Computer Science*, pages 232–242. Springer Berlin Heidelberg, 2007. 2
- [44] S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4):391–444, 2007. 1, 2, 5
- [45] S. Legg and J. Veness. An approximation of the universal intelligence measure. In *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*, pages 236–249. Springer, 2013. 3, 5, 12
- [46] L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973. 4
- [47] M. Li and P. Vitányi. *An introduction to Kolmogorov complexity and its applications (3rd ed.)*. Springer-Verlag, 2008. 1, 4
- [48] R. Madhavan, E. Tunstel, and E. Messina. *Performance Evaluation and Benchmarking of Intelligent Systems*. Springer, September, 2009. 2
- [49] A. Newell and H. A. Simon. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3):113–126, 1976. 2
- [50] J. Rothenberg, J. Paul, I. Kameny, J. R. Kipps, and M. Swenson. Evaluating expert system tools: A framework and methodology—workshops. Technical report, DTIC Document, 1987. 2
- [51] T. Schaul. An extensible description language for video games. *Computational Intelligence and AI in Games, IEEE Transactions on*, PP(99):1–1, 2014. 3
- [52] C. Schlenoff, H. Scott, and S. Balakirsky. Performance evaluation of intelligent systems at the national institute of standards and technology (nist). Technical report, DTIC Document, 2011. 2
- [53] H. A. Simon. Artificial intelligence: an empirical science. *Artificial Intelligence*, 77(1):95–127, 1995. 2
- [54] R. J. Solomonoff. A formal theory of inductive inference. Part I. *Information and control*, 7(1):1–22, 1964. 1
- [55] R. J. Sternberg (ed.). *Handbook of intelligence*. Cambridge University Press, 2000. 1

- [56] S. Whiteson, B. Tanner, M. E. Taylor, and P. Stone. Protecting against evaluation overfitting in empirical reinforcement learning. In *Adaptive Dynamic Programming And Reinforcement Learning (ADPRL), 2011 IEEE Symposium on*, pages 120–127. IEEE, 2011. [2](#)
- [57] R. Yonck. Toward a standard metric of machine intelligence. *World Future Review*, 4(2):61–70, 2012. [3](#)